Information Architecture for a Planetary "Exploration Web"

Norm Lamarra & Thom McVittie
Jet Propulsion Laboratory/Caltech
4800 Oak Grove, M/S 126-201
Pasadena, CA 91109
818-393-1561
Norm.Lamarra@jpl.nasa.gov
Thom.McVittie@jpl.nasa.gov

Abstract—The commercial internet is evolving service-oriented capabilities to simplify and enable new kinds of distributed applications in business, engineering, and management. This "web services" concept derives from a vision of the "connected economy", in which businesses recognize that their success in delivering their own customer services depends on efficient utilization of their providers' services in an increasingly inter-dependent economy. By "web services" we mean a common way of deploying distributed applications whose software components and data sources may be in different locations, formats, languages etc. Although such collaboration is not utilized significantly in planetary exploration, we believe there is significant benefit in developing an architecture in which missions could leverage each others capabilities. We believe that an incremental deployment of such an architecture could significantly contribute to the evolution of increasingly capable, efficient, and even autonomous remote exploration. The architecture thus supports a "global vision" of exploration, allowing each new mission to build upon others, not merely by reusing "heritage", but by planned IT infrastructure evolution. We are calling the resultant system a planetary "exploration web", and suggest that craft like rovers, landers, and orbiters should be enabled to collaboratively utilize all assets more effectively. We believe this will provide increased ROI, i.e., potentially much greater science return with lower operations cost, as well as making individual missions simpler and more robust, hence lower risk. Such an architecture would simplify access to information such as navigation, weather, terrain, and remote computation, and can become progressively capable as more services are deployed and more craft collaborate. As with the "web services" approach, participation can be made extremely simple (e.g., for low-cost "sensor web" micro-units), yet allowing resources such as bandwidth and storage can be more effectively shared (e.g., between a fixed micro-sensor, rover and orbiter). Further, this approach could assists in realizing the potential of agent-assisted exploration, whereby automated participants are both producing or consuming information at various levels.

TABLE OF CONTENTS

1. INTRODUCTION

Over the last few years, the "commercial" internet has evolved increasing capabilities to enable distributed and collaborative applications in business (like order processing and inventory control), in engineering (like automation from design through manufacturing) and in management (like human and knowledge resources). Today, the concept of "web services" has become a viable means to implement the "connected economy", rather than just a single "connected enterprise". Some businesses have recognized that their success in delivering customer services depends on efficient utilization of their providers' services, leading to a highly adaptive inter-dependent service economy. By "web services" we mean a common way of deploying distributed applications whose software components and data sources may be in different locations, formats, languages etc.; this was never truly achieved by any of the "distributed components" approaches promised several years ago, primarily because most required a homogeneous architecture (e.g., EJB) or platform (e.g., Windows). In fact, the current "web services" concept evolved partly in response to the

---

actual heterogeneity prevailing on the commercial web today.

2. VISION

The space IT component of planetary exploration in practice at JPL is far behind this vision, and in many cases has not even reached web capabilities from the early 1990's. Missions being proposed even for the 2007 timeframe still typically describe an isolated instrument suite perhaps with some data relay to "get the science home". By contrast, we present a much more revolutionary vision of IT-enabled planetary exploration, more akin to this decade's vision of the "web-based economy". In this vision, a new mission can easily participate in an evolving "exploration web", by leveraging existing capabilities and perhaps providing new ones. We believe that an incremental deployment of such an architecture could significantly contribute to the evolution of increasingly capable, efficient, and even autonomous remote exploration. The architecture thus supports a "global vision" of exploration, allowing each new mission to build upon others, not merely by reusing "heritage", but by planned IT infrastructure evolution.

The "business case" for such IT infrastructure is also not well understood in the space domain, except perhaps for standardized communication protocols. While these are indeed required, this is like saying: "the web is built on HTTP" (1990's terminology), rather than: "business effectiveness is built on adaptable web services" (2000-2010 terminology). The latter terminology derives from evolving application infrastructure (first CORBA and DCOM, then Java and EJB, then SOAP and XML).

We are calling the envisioned approach a planetary "exploration web", and suggest that craft like rovers, landers, and orbiters should be enabled to collaboratively utilize all assets more effectively. We believe this will provide increased ROI, i.e., potentially much greater science return with lower operations cost, as well as making individual missions simpler and more robust, hence lower risk. Such an architecture would simplify access to information such as navigation, weather, terrain, and remote computation, and can become progressively capable as more services are deployed and more craft collaborate.

The benefits of such a "global" perspective have also been recognized in the military arena. For example, in future battles, we expect our forces to prevail quickly with minimum losses, by being rapidly deployable, overwhelming in capability, and able to make decisions based on the "right" information (acquired from various sensors and delivered in consumable form to the "frontline warfighter") in real time. In this military domain, space infrastructure to enable this is currently being re-architected to foster "interoperability" not just between space and ground assets, but between military and civil infrastructures. Current investment in such development far exceeds that in the space exploration field, and is itself dwarfed by investment in the development of commercial web applications. Because of this investment (and JPL's involvement in part of this), we expect at least order-of-magnitude ROI in this effort, and possibly even two orders. It is also relevant to note that the Pentagon's long-held priorities on: 1 - platform (ship, plane), then 2 - sensor, then 3 - communications networks, are currently being reversed in order to recognize that they are in much greater need of better communication capability (which is relatively cheap) between existing sensors, rather than needing new sensors (which are expensive) or new platforms (which are even more expensive). This priority reversal also requires shifting the emphasis from stovepipe operations to flexible inter-operational and analysis scenarios, thus requiring more "commodity" information services (2010 timeframe and beyond). Our approach attempts to apply this thinking to planetary exploration.

3. OBJECTIVES

In the (highly-constrained) world of planetary exploration, technology infusion typically takes >5 years even when consciously planned for single technologies, and we predict much longer for infrastructure technologies (because typically several need to evolve together). In order to provide a significant increase in deployed capability for remote exploration in the 2010+ timeframe, we therefore propose sketching a "web service" space architecture this year. Recent history supports following this approach for the exploration arena: first, the successful "business web" was not architected "de jure" but grew "de facto" via evolving standards for interoperability as mentioned above; this means that many of the infrastructure "mistakes" (which we could not afford to make in "single" planetary missions) may have already been made (and surmounted) in the surviving successful internet. Second, the entire architecture doesn't need to be deployed at once, but can be progressively evolved with appropriate planning. The objective of our work is thus to make an intellectual connection between the web services architecture evolving in today's business internet, and the future "exploration web" described in the vision above. The specifics of what is currently being attempted are described in the Approach section below, but the overall effort is expected to be multi-year and progressive.

4. APPROACH

In order to support this vision, it is necessary (but not sufficient) to build interoperability between mission assets at three levels: application, middleware, and space protocol. Applications (software modules, agents, and mission components) must be enabled to share and process data and information seamlessly even when these parts have been implemented (and perhaps designed) independently. We therefore propose to apply a "web services" approach to the capture, location and exchange of data products between spaceborne assets participating in the "exploration web".

We would adapt this approach to enable autonomous systems (in space) to assign meaning to sensor interfaces and derive information from data (initially in predetermined ways). This service approach can also be developed for ground systems in the same way (and in fact could be used as a testbed for future space deployment).

Eventually, a researcher or analyst may be able to ask a high-level question, which would be progressively broken down into many (automated) tasks involving data capture, processing, analysis, and transport. Space and ground components would thus support the architecture, data services, and algorithms for intelligent data access and retrieval, but would be tailored for each environment (space vice ground), and built upon common standards in each domain (e.g., CCSDS protocols vice IP protocols).

Recognizing the terrestrial evolution towards "loosely-coupled" services, we believe that space interoperability can be built most flexibly and simply upon message-oriented communications. We are currently leveraging an in-house investment in such robust messaging middleware: JPL has developed SharedNet for the U.S. Marines (Sea Dragon), to allow dispersed units to communicate information asynchronously over heterogeneous and intermittent networks, including low-bandwidth radios. In an FY02 prototype, we exploited the obvious similarities in both the constraints and goals between this harsh military environment and the space exploration environment. We selected an autonomous planning scenario, in which a conflict between rover and orbiter plans (e.g., MER and Odyssey) was resolved for a case where the rover could not obtain a direct Earth connection quickly enough. The resulting plan, negotiated between the rover and orbiter, achieved improved science operation despite the constrained Earth connectivity.

We are thus trying to generalize "remote communication" into a high-level architecture for cooperation among space assets. First, we are attempting to define high-level data objects and a mechanism to allow efficient exchange of such objects between assets (producers and consumers). This also allows on-board applications to be insulated from inessential details, such as the vagaries of the space communication. We also need to provide interfaces to appropriate higher-level peer information services, and to design these services to be layered appropriately given the highly-constrained "network" environment and resource availability. We are currently defining a prototype combining these elements to determine feasibility and benefits. This approach thus mimics the development of terrestrial "web services" based on encapsulated components communicating via Internet protocols.

We distinguish the proposed "exploration web" architecture from the (required) underlying communication protocols, just as the "business web services" architecture is distinguished from the underlying internet protocols (e.g., HTTP and TCP). However, we cannot develop or apply such an architecture for space without careful attention to the constraints (e.g., intermittent connectivity, progressive deployment, limited resources such as bandwidth, power, etc.). Our approach would become more obviously beneficial as more services are deployed and more craft "co-operate". However, without that vision, the required infrastructure will never be developed.

5. STRATEGIC BENEFITS

The proposed information architecture could provide strategic benefits in several areas, as shown in the following examples.

Autonomous mobility: During rover traverse, information such as high-resolution imagery could help plan a path avoiding obstacles. If such information could be obtained directly from local sensors, a round trip to Earth would be avoided, thus improving decision response-time as well as assisting opportunistic science. Such local access requires the ability to locate and access stored data, both elements provided by our proposed information architecture. Similarly, weather, navigation, or even computation services (like planning) can be supplied locally more efficiently than via Earth.

Deep Space Communication: Standardized protocols (e.g., CCSDS) can help optimize use of constrained resources such as bandwidth, power, etc. However, similar standardization at the application layer can improve use of the communication channel itself, by allowing standard ways for applications to manage and transport data, even when not in communication. For example, a rover could verify its own position observations by receiving mapping information from an orbiting mapping sensor.

Formation flying: Our approach significantly simplifies the local distribution of information (eg, location, time, sensor data). Collaborative planning can also be significantly simplified as demonstrated by our FY02 prototype application built using SharedNet communications.

Advanced spacecraft computing and autonomy: Our architecture simplifies access to remote computing resources, thus allowing a simple, low-cost node to perform more complex work via leveraging of off-board assets. Autonomy attempts to make decisions based on the available information; our exploration web could provide more timely and useful information in appropriate form (e.g., from a remote sensor feed) to assist such decision-making. As mentioned above, this approach can also dramatically simplify spacecraft "collaboration" (e.g., distributed planning). From a goal-based (science or engineering) perspective, our approach supports distribution and negotiation of lower-level activities (elaborated from high-level goals) among cooperating (communicating) craft.

3

Active and passive remote sensors: A sensor web can be a simple version of the exploration web suggested here. Features like enhanced data collection, distribution and processing can easily be added for all types of sensor information at various levels (from raw to interpreted), thus effectively leveraging the potential contribution of every sensor. A passive sensor could simply "publish" regular measurements for consumption and action by more complex active nodes.

Agent communication: Agents process known types of information by applying rules to create actionable outputs. A simple example is an on-board agent reporting vehicle health or sensor status to external recipients. More complex higher-level agents could then provide higher-level processed information, e.g., generating a science alert based on receiving and correlating distributed sensor data.

6. CONCLUSIONS

We have embarked on building an information architecture which can assist remote spacecraft in participating in JPL's proposed "planetary internet". This is a pilot activity based on the evolving "web services" approach, and is intended to be incremental and layered. We are currently prototyping a distributed science application to demonstrate simplified data management and processing (example higher-layer services).

Norm Lamarra is a software developer in the Engineering and Communications Infrastructure Section of JPL's Information Systems Division. He holds B.Sc. and M.Sc. degrees from the University of Birmingham, UK, in Mathematical Physics and Electronic Engineering respectively, and a Ph.D. from UCLA in System Science (Communication Systems).